Experimental noise filtering by quantum control

A. Soare,[1,*] H. Ball,[1,*] D. Hayes,[1,†] J. Sastrawan,[1] M. C. Jarratt,[1]
J. J. McLoughlin,[1] X. Zhen,[2] T. J. Green,[1] and M. J. Biercuk[1,‡]

[1] *ARC Centre for Engineered Quantum Systems, School of Physics, The University of Sydney, NSW 2006 Australia*
National Measurement Institute, West Lindfield, NSW 2070 Australia
[2] *Tsinghua University, Beijing, People's Republic of China*

(Dated: July 8, 2014)

Extrinsic interference is routinely faced in systems engineering, and a common solution is to rely on a broad class of *filtering* techniques in order to afford stability to intrinsically unstable systems or isolate particular signals from a noisy background. For instance, electronic systems are frequently designed to incorporate electrical filters composed of, *e.g.* RLC components [1], in order to suppress the effects of out-of-band fluctuations that interfere with desired performance. Experimentalists leading the development of a new generation of quantum enabled technologies similarly encounter time-varying noise in realistic laboratory settings. They face substantial challenges in either suppressing such noise for high-fidelity quantum operations [2–6] or controllably exploiting it in quantum-enhanced sensing [7–11] or system identification tasks [12], due to a lack of efficient, validated approaches to understanding quantum dynamics in the presence of time-varying noise. In this work we use the theory of quantum control engineering [13, 14] and experiments with trapped $^{171}\mathrm{Yb}^{+}$ ions to construct novel *noise spectral filters* with user-defined properties and compatible with incorporation into arbitrary control operations. Our results provide the first experimental validation of generalized filter-transfer functions for arbitrary quantum control operations [15, 16], and demonstrate their utility for developing novel robust control and sensing protocols. We provide a detailed framework for filter synthesis appropriate for arbitrary single-qubit state transformations and experimentally validate the performance of the resulting noise filters. These experiments provide a significant advance in the theory of quantum control and unlock new capabilities for the emerging field of quantum systems engineering.

arXiv:1404.0820v2 [quant-ph] 5 Jul 2014

The presence of time-dependent noise in either the environment or control of a single qubit randomizes the intended evolution trajectory of a state $|\psi_0\rangle \rightarrow |\psi_\mathrm{T}\rangle$ (Fig. 1a), ultimately reducing the fidelity of the operation. This is captured as $\mathcal{F}_\chi(\tau) = \frac{1}{2}(1 + e^{-\chi(\tau)})$, where $\chi(\tau) = \frac{1}{\pi}\sum_i \int_0^\infty d\omega S_i(\omega) F_i(\omega)$, and τ is the total duration of the operation. In this expression for fidelity, the integral considers contributions from independent noise processes through their frequency-domain power spectra $S_i(\omega)$, $i \in \{z, \Omega\}$, capturing dephasing along $\hat{z}$ and amplitude noise co-rotating with a resonant drive field (see *Supplementary Material*).

The quantities $F_i(\omega)$ describe the effective frequency response of any applied control operation (including free evolution), and are referred to as filter transfer functions, "FF"s for the control, providing an efficient manner to understand the dynamical response of a controlled quantum system in a time-varying environment [17–22]. Despite their appeal – FFs may be characterized using a standard engineering approach considering frequency passbands, stopbands, and filter order – to date, the only demonstrations of FFs in quantum control have been undertaken in the simple case of application of the *identity* operator (dynamical decoupling) [12, 23, 24]. In this case the analytics are dramatically simplified due to assumptions of unbounded "bang-bang" control.

More generally, calculating $F_i(\omega)$ and hence characterizing the spectral response of an *arbitrary* bounded-strength control operation is challenging in cases where the noise and control Hamiltonians do not commute (e.g. a driven σ_x rotation in the presence of σ_z dephasing). With recent theoretical developments FFs may now be calculated analytically for arbitrary control [15, 16]; it is this more general case where the impact of noise filtering and the filter transfer functions may have the most significant impact on the quantum engineering community, and where experimental tests are most germane.

To see this we may consider the various tasks that might be of interest to an experimentalist engaged in quantum engineering and the role of noise spectral filtering in these applications. Noise filtering itself is achieved through construction of a control protocol (Fig. 1a) which reduces the *controllability* of the quantum system by the noisy environment *over a defined frequency band* by suitably modifying $F_i(\omega)$. In quantum information an experimentalist may aim to suppress broadband low-frequency noise in order to maximize the fidelity of a bounded-strength quantum logic operation (Fig. 1b, upper trace). Alternatively, in quantum enabled sensing or system identification he or she may perform narrowband spectral characterization of a given operation (Fig. 1b, lower trace), where the measured infidelity under filter application represents the signal of interest [11, 12]. The FFs are simple analytic objects which enable the extraction of exactly this information calculated for arbitrary control and arbitrary universal noise.

In our experimental system, based on the 12.6 GHz qubit transition in $^{171}\mathrm{Yb}^{+}$ (see *Supplementary Material*),

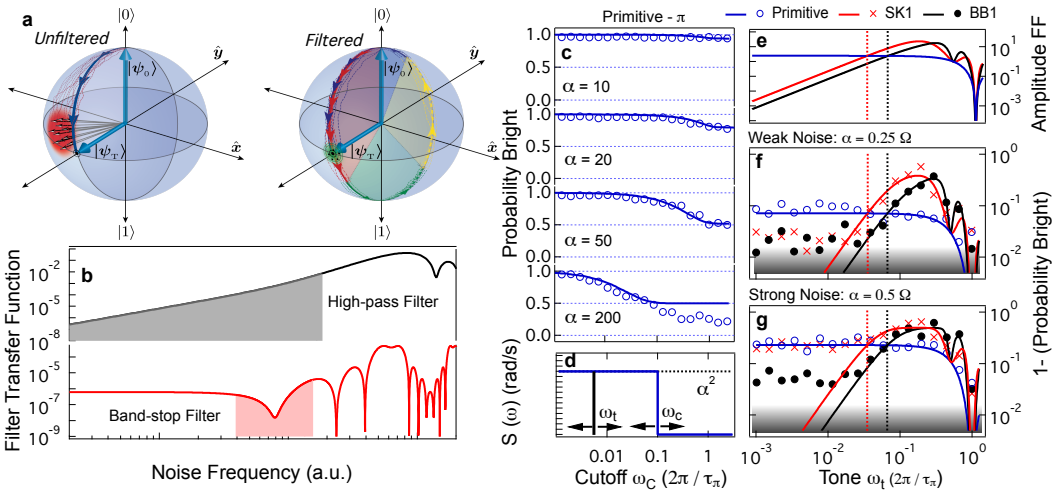

FIG. 1. Noise filters and experimental validation of the predictive power of the filter transfer function. a) Time-varying noise during an operation (a rotation on the Bloch sphere, here $\theta = \pi/2$, $\phi = 0$) produces a broad range of outcomes (red uncertainty cone, left) and may yield an offset of the average final state from the target state, measured as operational infidelity. Filtered trajectory, depicted as a user-defined modulation pattern on the control (colored segments), changes the measured fidelity by reducing the uncertainty due to noise in a specified band. b) Schematic representation of noise filters of interest - shaded areas represent filter stop-bands - crafted by control modulation as indicated above. c) Measurements of operational fidelity with engineered dephasing noise for primitive π rotation, $|0\rangle \rightarrow |1\rangle$, as a function of dimensionless noise cutoff frequency overlaid with FF-based calculations of $\mathcal{F}_\chi(\tau)$. Decay to value 0.5 corresponds to full decoherence. For $\omega_c/2\pi = 1$ the highest frequency contribution to $S_z(\omega)$ undergoes a complete cycle of oscillation over τ_π, indicating that the noise is time-dependent on the scale of a single experiment even for $\omega_c/2\pi \ll 1$. Each data point is the result of averaging over 50 different noise realizations. (d) Schematic representation of the quasi-white noise power spectrum employed in (c) and single-tone power spectrum employed in (f-g). Noise strength parameterized by α. e) Calculated $F_\Omega(\omega)$, for primitive and compensating π pulses (see [25]). Vertical lines indicate frequencies where $F_\Omega(\omega)$ for SK1 (red) and BB1 (black) cross primitive (blue), indicating an expected inversion of performance. f)-g) Swept-tone multiplicative amplitude noise measurements , $S_\Omega(\omega) \propto \delta(\omega_t - \omega)$, for various π rotations (averaged over 20 noise realizations). Vertical axis is a proxy for measured operational *infidelity*. Solid lines indicate $1 - \mathcal{F}_\chi(\tau)$, revealing good agreement in the weak noise limit (f) across three decades of frequency, down to measurement fidelity limit, $\sim 98.5\%$, indicated by grey shading. Measured gate-error crossover points correspond well with crossovers in the FFs for these gates (vertical dashed lines). Detailed performance differences between protocols in the low-error limit can be revealed through randomized benchmarking, as performed later (Fig. 3). g) Strong-error limit, first-order approximations are violated and contributions from higher-order Magnus terms contribute to the measured error in the low-frequency limit, yielding (expected) differences between SK1 and BB1 not captured by the FF.

we are able to perform quantitative tests of operational fidelity revealing the spectral characteristics of arbitrary control operations; these may then be compared against calculations of $\mathcal{F}_\chi(\tau)$ as a fundamental test of FF validity. A key tool in our studies is bath engineering [26], in which we add noise with user-defined spectral characteristics to the control system, producing well controlled unitary dephasing or depolarization.

In our first experiment we measure the probability that a π_x-pulse drives qubit population from the dark state to the bright state, $|0\rangle \rightarrow |1\rangle$, while varying the high-frequency cutoff, ω_c, of an engineered non-Markovian dephasing bath (Fig. 1c). As the high-frequency cutoff of the noise is increased and fluctuations fast compared to the control (τ_π) are added to the noise power spectrum, $S_z(\omega)$, errors accumulate reducing the measured fidelity. We calculate $\mathcal{F}_\chi(\tau)$ using the form of the noise and the analytic FF for a driven primitive gate under dephasing [16], finding good agreement with experimental measurements using *no free parameters*, thus verifying the predictive power of the FF.

The FFs for much more complex control such as compensating composite pulses [25, 27] can be calculated as well [28], revealing their sensitivity to time-dependent noise - an important characteristic for deployment in quantum information settings. We experimentally demonstrate a form of quantum system identification, reconstructing the amplitude-noise filter functions, $F_\Omega(\omega)$, for SK1 and BB1 π_x-pulse sequences by measuring operational fidelity in the presence of a swept narrow-band amplitude noise signal [26]. The resulting measured (in)fidelity effectively traces out the FF; key features such as performance-crossover frequencies between primitive and compensating gates and deep notches in the filter at high frequency are quantitatively reproduced in experimental measurements. Again, first-order fidelity calculations, $\mathcal{F}_\chi(\tau)$, match data well in the weak noise limit (Fig. 1f) with *no free parameters*.

Ultimately, the underlying physical principles giving rise to the analytic form of $F_i(\omega)$ are based on the well tested average Hamiltonian theory [29, 30] common to NMR and quantum information. Despite this shared theoretical foundation, the calculation of spectral filtering properties is quite distinct from finding compensating-pulse protocols giving error-compensation in a Magnus expansion. This distinction is important as time-varying colored classical noise is commonly encountered in laboratory settings, but high-Magnus-order composite pulses need not be efficient noise spectral filters (see *Supplementary Material* and [28]).

Measurements of SK1 and BB1 fidelity highlight this difference. Both gates provide similar filtering of time-dependent noise, given by the *filter order*, Fig. 1e, despite the significant differences in their construction; BB1 is designed to suppress higher-order Magnus terms than SK1 (see *Supplementary Material*). In the strong-error limit (Fig. 1g), however, significant performance deviations arise at low frequencies; the increased Magnus-order cancellation of BB1 improves quasi-static error suppression relative to SK1 (the regime for which both sequences were originally crafted). Overall, frequency-domain characteristics are captured accurately through the FF in the weak noise limit (Fig. 1f), but would not be directly evident through a Magnus expansion.

These simple but powerful validations of the generalized filter transfer function's predictive power now open the possibility of demonstrating the *construction* of noise filters with a specified spectral response, employing the filter transfer functions as key analytic tools. Filters may take a wide variety of forms - including high-pass filters for broadband noise suppression and band-stop filters useful for narrowband noise characterization (Fig. 1b).

In the example that follows, we focus on a common setting in which we aim to improve operational fidelity by reducing the influence of broadband non-Markovian noise on a target state transformation. Filters are realized as sequences of time-domain control operations with tunable pulse amplitude and phase, similar in spirit to compensating composite pulses in NMR [25, 27, 28], dynamically corrected gates (DCGs) in quantum information [31, 32], and open-loop modulated pulses in quantum control [33, 34]. In this setting we wish to synthesize a filter with arbitrary, user-defined spectral characteristics captured by a cost-function, $A(\mathbf{\Gamma}_n)$, to be minimized over n pulse segments in a filter construction.

With the unique task of creating filters tailored to a given noise spectrum in hand, we introduce a basic framework for filter construction leveraging the filter-transfer function. An arbitrary n-segment filter is represented over successive timesteps through the matrix quantity $\mathbf{\Gamma}_n(\theta_l, \tau_l, \phi_l)$ (Fig. 2b, *Supplementary Material*) describing the properties of a near-resonant carrier frequency enacting driven operations. In each segment of duration τ_l we perform a driven operation generating a rotation through an angle $\theta_l = \int_{t_{l-1}}^{t_l} \Omega_l(t)dt$ about the axis $\vec{r}_l = (\cos(\phi_l), \sin(\phi_l), 0)$, with $\Omega_l(t)$ the Rabi rate over the lth pulse segment.

To provide *efficient* solutions to filter design we restrict our control space and focus on constructions synthesized using concepts from functional analysis in the basis set of *Walsh functions* - square-wave analogues of

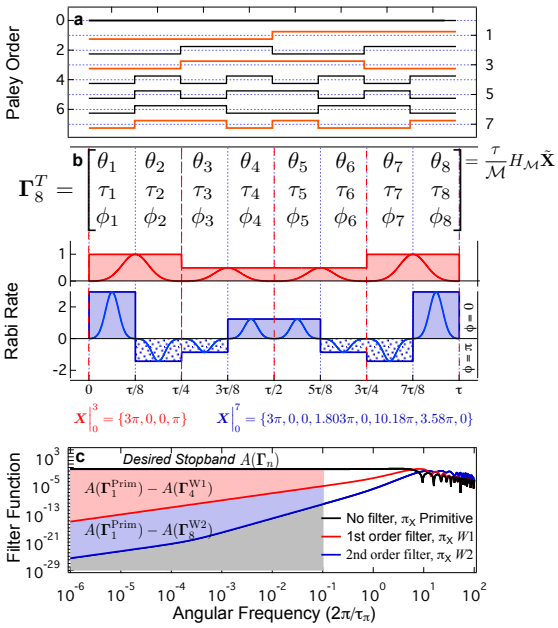

FIG. 2. Synthesis of high-pass amplitude-modulated filters from the Walsh functions. a) The first eight Walsh functions used in filter synthesis, $\{\mathrm{PAL}_0, \mathrm{PAL}_7\}$, with maximum-Hamming-weight-indexed functions highlighted. b) Representative amplitude profiles for filter constructions found via numerical search over the Walsh basis with four (red, denoted $W1$) and eight (blue, denoted $W2$) time steps. Vertical axis represents Ω, the Rabi rate per time step; negative values indicate π-phase shifts. Synthesis may be performed over square (flat-top) pulse-segments or Gaussian-shaped pulse segments with results differing only in the resulting Walsh coefficients. The matrix representing filter characteristics over eight segments is superimposed on the amplitude profiles (for $n = 4$, neighboring segments between red dashed lines are combined). The first row (the angles of rotation in each segment of the filter) is determined via Walsh synthesis, indicated by the vectors $\boldsymbol{X}^T\big|_0^n$, containing the spectral weights over $\mathrm{PAL}_0 \rightarrow \mathrm{PAL}_n$. In the case of Gaussian pulse envelopes Walsh synthesis sets the first line, θ_l. The symbol $\tilde{\boldsymbol{X}}$ indicates reordering for Hadamard synthesis, with listed coefficients appropriate for square pulse envelopes. c) The filter transfer function for a primitive π_x rotation and for synthesized noise filters. Performance improvement over the desired stopband of the filter captured in cost function $A(\mathbf{\Gamma}_{4(8)}^{W1(W2)})$ and its difference relative to that for the primitive operation, $A(\mathbf{\Gamma}_1^{\mathrm{Prim}})$. Filter $W1$ gives improvement indicated by the red shading, with additional improvement in the cost function given by $W2$ indicated by blue shading.

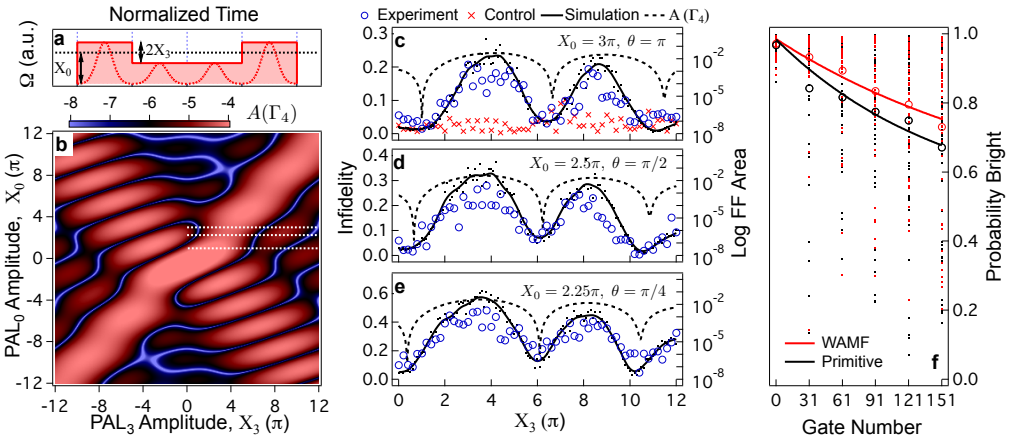

FIG. 3. Construction of the first-order Walsh amplitude modulated dephasing-suppressing filter. a) Schematic representation of Walsh synthesis for a four-segment amplitude-modulated filter (WAMF). Walsh synthesis can be used to determine either the modulating envelope of square pulse segments, or the net area of discrete Gaussian pulses with differing amplitudes. b) Two-dimensional representation of the integral metric defining our target cost function, $A(\mathbf{\Gamma}_4)$ integrated over the stopband $\omega \in [10^{-9}, 10^{-1}]\tau^{-1}$. Areas in blue minimize $A(\mathbf{\Gamma}_4)$, representing effective filter constructions. The X_0 determines the net rotation enacted in a gate while X_3 determines the modulation depth, as represented in a). White lines indicate possible constructions for filters implementing rotations of $\theta = \pi$, $\theta = \pi/2$, and $\theta = \pi/4$ from top to bottom. c)-e) Experimental measurement of gate infidelity (left axis) for rotations constructed from various Walsh coefficients in the presence of engineered noise ($\omega_c/2\pi = 20$ Hz). Black dots and line represented calculated fidelity by Schroedinger equation integration (raw and smoothed respectively). All values of X_3 for a given X_0 implement the same net rotation, indicated by control experiment with no noise. Total rotation time is scaled with X_3 to preserve a maximum *Rabi rate*. Black dashed line (right axis) corresponds to $A(\mathbf{\Gamma}_4)$ from panel (a). In experiments we always perform a net π rotation $|0\rangle \rightarrow |1\rangle$ by sequentially performing identical copies of rotations for $\theta < \pi$. f) Randomized benchmarking results (50 randomizations) demonstrating superior performance of modulated gate in the small-error limit (infidelity $< 0.5\%$ per gate), see *Supplementary Material*.

the sines and cosines [11, 35, 36] (Fig. 2a). This approach is by no means the only basis set for composite gate construction [37, 38], but provides significant benefits for our problem [36]. For instance, their piecewise-constant construction builds intrinsic compatibility with discrete clocking [39] and classical digital logic, while the well characterized mathematical properties of the Walsh functions provide a basis for establishing simple *analytic* filter-design rules, and flexibility in realizing a wide variety of filter forms (see *Supplementary Material*).

As an example we synthesize filters via weighted linear combination of Paley-ordered Walsh functions, $\mathrm{PAL}_k(x)$, designed to suppress time-varying dephasing noise over a low-frequency stopband while implementing a bounded-strength driven rotation about the x-axis on the Bloch sphere. In this case the Walsh-synthesized waveform dictates an amplitude modulation pattern for the control field over discrete time-segments. Importantly, Walsh filter synthesis is compatible with pulse segments possessing *arbitrary pulse envelope*, including sequences of *e.g.* square (used here) or Gaussian pulse segments (Fig. 2b).

Walsh-synthesis design rules dictate that we implement our filtered rotation by θ_x over a minimum of four discrete steps, permitting synthesis over PAL_0 to PAL_3. Within this small set, the coefficient of PAL_0, denoted X_0, sets the total rotation angle $\theta \mod 2\pi$ for the modulated driven evolution, and only nonzero X_3 preserves symmetry. Relative to an unfiltered primitive gate, $F_z(\omega)$ for a construction giving effective filtered evolution, $W1$, shows increased steepness in the stopband (Fig. 2c, red), reducing $A(\mathbf{\Gamma}_4)$ (here, $\theta = \pi$). Doubling the number of pulse-segments, and allowing Walsh synthesis over coefficients $\{X_0, ..., X_7\}$ yields filter $W2$, which increases the order of time-domain noise filtering (steeper rolloff), further reducing the cost function for optimization (blue shaded area in Fig. 2c). Interestingly, $W1$ is a special case of an analytically constructed dynamically corrected NOT gate (a π-rotation) [31].

Filters $W1$ and $W2$ were found using a Nelder-Mead simplex optimization over Walsh coefficients and are representative, rather than unique solutions. In Fig. 3b we show the calculated cost function, $A(\mathbf{\Gamma}_4)$, as a function of X_0 and X_3. Blue areas meet our minimized target indicating useful filters, revealing a wide variety of possible constructions with favorable filtering characteristics.

We experimentally test the performance of four-segment amplitude-modulated filters by scanning over X_3 for fixed X_0 (denoted by white dotted lines in Fig. 3b). Values of X_3 minimizing $A(\mathbf{\Gamma}_4)$ (dips in the dashed trace, right axis) also minimize the experimentally measured infidelity in the presence of engineered low-frequency noise (open circles, left axis). This behav-

ior is observed for various target rotation angles of interest (Fig. 3c-e), with predicted shifts in the optimal values of X_3 with changes in X_0 borne out through experiment. In the small error limit we also show that a filtered π pulse ($W1$) outperforms the best primitive operation via randomized benchmarking (Fig. 3f), despite being four times longer. These measurements thus validate the concept of spectral filter synthesis based on analytic FFs.

In summary, our experiments have shown that filter transfer functions are effective tools for the characterization of complex time-domain control operations in the presence of universal noise, revealing fine details about their spectral features. We have provided the first experimental validation of their predictive power for arbitrary quantum logic, observing good agreement between measurements on trapped ions and FF-calculated fidelity using *no free parameters*. In addition, we have shown how one may leverage these analytic constructs for the synthesis of a wide range of time-domain noise filters with user-specified characteristics from the Walsh basis functions.

Our focus has been on providing a validated framework for the vital task of filter construction to complement existing techniques, rather than attempting to find maximum-performance error-robust gates. Our results on high-pass noise filters, for instance, add to existing compensating pulse sequences designed for quasi-static noise, as well as gate constructions with interleaved dynamical decoupling that seek to periodically "refocus" a quantum trajectory [40–44]. Other applications of this approach include crafting notch filters for noise characterization tasks and extension to the filtering of noise in multiqubit logical operations [45]. We believe that this simple framework will provide a straightforward path for experimentalists to characterize and suppress the effects of noise in generic quantum coherent systems, ultimately enabling a new generation of engineered quantum technologies.

We thank P. Fisk, M. Lawn, M. Wouters, and B. Warrington for technical assistance and K. Brown, J.T. Merrill and L. Viola for useful discussions. This work partially supported by the US Army Research Office under Contract Number W911NF-11-1-0068, and the Australian Research Council Centre of Excellence for Engineered Quantum Systems CE110001013, the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), through the Army Research Office, and the Lockheed Martin Corporation.

* *These two authors contributed equally to this work.*
† *Present address:* Lockheed Martin Corporation
‡ michael.biercuk@sydney.edu.au

[1] D. Pozar, *Microwave Engineering, 4th ed.* (John Wiley & Sons, Hoboken, NJ, 2011).
[2] P. O. Schmidt, T. Rosenband, C. Langer, W. M. Itano, J. C. Bergquist, and D. J. Wineland, Science. **309**, 749 (2005).
[3] S. Foletti, H. Bluhm, D. Mahalu, V. Umansky, and A. Yacoby, Nature Physics (2009).
[4] J. W. Britton, B. C. Sawyer, A. C. Keith, C.-C. J. Wang, J. K. Freericks, H. Uys, M. J. Biercuk, and J. J. Bollinger, Nature **484**, 489 (2012).
[5] A. Smith, B. E. Anderson, H. Sosa-Martinez, C. A. Riofrio, I. H. Deutsch, and P. S. Jessen, Phys. Rev. Lett. **111**, 170502 (2013).
[6] R. Islam, C. Senko, W. C. Campbell, S. Korenblit, J. Smith, A. Lee, E. E. Edwards, C.-C. J. Wang, J. K. Freericks, and C. Monroe, Science **340**, 583 (2013).
[7] M. S. Grinolds, S. Hong, P. Maletinsky, L. Luan, M. D. Lukin, R. L. Walsworth, and A. Yacoby, Nature Phys. **9**, 215 (2013).
[8] F. Dolde, H. Fedder, M. W. Doherty, T. Nbauer, F. Rempp, G. Balasubramanian, T. Wolf, F. Reinhard, L. C. L. Hollenberg, F. Jelezko, and J. Wrachtrup, Nature Phys. **7**, 459 (2011).
[9] G. Kucsko, P. C. Maurer, N. Y. Yao, M. Kubo, H. J. Noh, P. K. Lo, H. Park, and M. D. Lukin, Nature **500**, 54 (2013).
[10] F. Shi, X. Kong, P. Wang, F. Kong, N. Zhao, R.-B. Liu, and J. Du, Nat. Phys. **10**, 21 (2014).
[11] A. Cooper, E. Magesan, H. Yum, and P. Cappellaro, Nature Comms. **5**, 3141 (2014).
[12] J. Bylander, S. Gustavsson, F. Yan, F. Yoshihara, K. Harrabi, G. Fitch, D. G. Cory, Y. Nakamura, J.-S. Tsai, and W. D. Oliver, Nat. Phys. **7**, 565 (2011).
[13] J. W. Clark, D. G. Lucarelli, and T. J. Tarn, Int. J. Mod. Phys. B **17**, 5397 (2003).
[14] L. Bouten, R. V. Handel, and M. R. James, SIAM J. Control Optim. **46**, 2199 (2007).
[15] T. J. Green, H. Uys, and M. J. Biercuk, Phys. Rev. Lett. **109**, 020501 (2012).
[16] T. J. Green, J. Sastrawan, H. Uys, and M. J. Biercuk, New J. Phys. **15**, 095004 (2013).
[17] We employ here the so-called modified filter transfer function, which subsumes a factor of ω^{-2} into the definition of $F_i(\omega)$. See [12, 22] for details.
[18] A. Kofman and G. Kurizki, Phys. Rev. Lett. **87**, 270405 (2001).
[19] A. G. Kofman and G. Kurizki, Phys. Rev. Lett. **93**, 130406 (2004).
[20] G. Uhrig, Phys. Rev. Lett. **98**, 100504 (2007).
[21] L. Cywinski, R. M. Lutchyn, C. P. Nave, and S. D. Sarma, Phys. Rev. B **77**, 174509 (2008).
[22] M. J. Biercuk, A. C. Doherty, and H. Uys, J. Phys. B **44**, 154002 (2011).
[23] M. J. Biercuk, H. Uys, A. P. VanDevender, N. Shiga, W. M. Itano, and J. J. Bollinger, Nature **458**, 996 (2009).
[24] G. A. Álvarez and D. Suter, Phys. Rev. Lett. **107**, 230501 (2011).
[25] J. T. Merrill and K. R. Brown, arXiv:1203.6392 (2012).
[26] A. Soare, H. Ball, D. Hayes, X. Zhen, M. C. Jarratt, H. Uys, and M. J. Biercuk, arXiv (2014), 1402.4632.
[27] L. M. K. Vandersypen and I. Chuang, Rev. Mod. Phys. **76**, 1037 (2004).
[28] C. Kabytayev, T. J. Green, K. Khodjasteh, M. J.

Biercuk, L. Viola, and K. R. Brown, arXiv (2014), 1402.5147.
[29] W. Magnus, Commun. Pure Appl. Math. **7**, 649 (1954).
[30] U. Haeberlen and J. S. Waugh, Phys. Rev **175**, 453 (1968).
[31] K. Khodjasteh and L. Viola, Phys. Rev. Lett. **102**, 80501 (2009).
[32] X. Wang, L. Bishop, J. Kestner, E. Barnes, K. Sun, and S. D. Sarma, Nature Comms. **3**, 997 (2012).
[33] B. Fauseweh, S. Pasini, and G. Uhrig, Phys. Rev. A **85**, 022310 (2012).
[34] F. Ticozzi, K. Nishio, and C. Altafini, IEEE Trans. Auto. Control **58**, 74 (2013).
[35] K. G. Beauchamp, *Walsh Functions and their Applications* (Academic Press, London, 1975).
[36] D. Hayes, K. Khodjasteh, L. Viola, and M. J. Biercuk, Phys. Rev. A **84**, 062323 (2011).
[37] P. Owrutsky and N. Khaneja, arXiv:1204.006 (2012).
[38] N. C. Jones, T. D. Ladd, and B. H. Fong, New J. Phys. **14**, 093045 (2012).
[39] T. E. Hodgson, L. Viola, and I. D'Amico, Phys. Rev. A **81**, 062321 (2010).
[40] C. Barthel, J. Medford, C. M. Marcus, M. P. Hanson, and A. C. Gossard, Phys. Rev. Lett. **105**, 266808 (2010).
[41] M. D. Shulman, O. E. Dial, S. P. Harvey, H. Bluhm, V. Umansky, and A. Yacoby, Science **336**, 202 (2012).
[42] G.-Q. Liu, H. C. Po, J. Du, R. B. Liu, and X. Y. Pan, Nat. Comms. **4**, 2254 (2013).
[43] A. M. Souza, G. A. Álvarez, and D. Suter, Phys. Rev. A **86**, 050301 (2012).
[44] T. van der Sar, Z. Wang, M. Blok, H. Bernien, T. Taminiau, D. Toyli, D. Lidar, D. Awschalom, R. Hanson, and V. Dobrovitski, Nature **484**, 82 (2012).
[45] D. Hayes, S. M. Clark, S. Debnath, D. Hucul, I. V. Inlek, K. W. Lee, Q. Quraishi, and C. Monroe, Phys. Rev. Lett. **109**, 020503 (2012).
[46] K. J. Horadam, *Hadamard Matrices and their Applications* (Princeton University Press, Princeton, 2007).
[47] D. A. Lidar and T. A. Brun, *Quantum Error Correction* (Cambridge University Press, New York, 2013).
[48] S. Olmschenk, K. C. Younge, D. L. Moehring, D. N. Matsukevich, P. Maunz, and C. Monroe, Phys. Rev. A **76**, 052314 (2007).

Supplementary Material

We consider a model quantum system consisting of an ensemble of identically prepared noninteracting qubits immersed in a weakly interacting noise bath and driven by an external control device. Working in the interaction picture with respect to the qubit splitting ω_a state transformations are represented as unitary rotations of the Bloch vector. The generalized time-dependent Hamiltonian is then written

$$H(t) = H_c(t) + H_0(t) \tag{1}$$

where $H_c(t)$ describes perfect control of the qubit state, e.g. via an ideal external driving field, and the noise Hamiltonian $H_0(t)$ captures undesirable interactions with a (universal) noise bath.

The specific forms taken by $H_c(t)$ and $H_0(t)$ in this work are given in the sections below, where we treat both dephasing (detuning) and amplitude-damping (coherent relaxation) noise processes. We will begin with this model to craft time-dependent noise filters, and detail this method in the following sections.

Defining the control space

Representing the qubit state on the Bloch sphere, state manipulation maps to a rotation of the Bloch vector in $\mathbb{R}^3$ and described by the unitary $U(\theta, \hat{\sigma}_{\vec{r}}) := \exp\left(\frac{-i\boldsymbol{\sigma}\cdot\vec{\boldsymbol{r}}\theta}{2}\right)$, reflecting the homeomorphism between $SU(2)$ and $SO(3)$. In effect, the spin operator $\hat{\sigma}_{\vec{r}} := \vec{\boldsymbol{r}}\cdot\boldsymbol{\sigma}$ generates a rotation though an angle θ about an axis defined by the unit vector $\vec{\boldsymbol{r}} \in \mathbb{R}^3$. For our purpose control takes the form of a composite pulse sequence consisting of n such unitaries executed over a time period $[0, \tau]$, with the lth pulse in the sequence written

$$P_l := U(\theta_l, \hat{\sigma}_{\phi_l}) = \exp\Big[-i\frac{\hat{\sigma}_{\phi_l}}{2}\int_{t_{l-1}}^{t_l}\Omega_l(t)dt\Big] \tag{2}$$

$$\hat{\sigma}_{\phi_l} := \cos(\phi_l)\hat{\sigma}_x + \sin(\phi_l)\hat{\sigma}_y. \tag{3}$$

Here $\Omega_l(t)$ is the Rabi rate with arbitrary amplitude envelope in a single pulse, $\tau_l = t_l - t_{l-1}$ is the pulse duration, and the spin operator $\hat{\sigma}_{\phi_l}$, parametrized by $\phi_l \in [0, 2\pi]$, generates a rotation $\theta_l = \int_{t_{l-1}}^{t_l}\Omega_l(t)dt$ of the Bloch vector about an axis $\vec{\boldsymbol{r}}_l \equiv (\cos(\phi_l), \sin(\phi_l), 0)$ in the xy-plane[1]. This sequence of control unitaries implies a natural partition of the total sequence duration τ into n subintervals $I_l = [t_{l-1}, t_l]$, $l \in \{1, n\}$, such that the lth pulse has duration $\tau_l = t_l - t_{l-1}$ with t_{l-1} and t_l the start

[1] For a resonantly driven qubit ϕ_l is the phase of the driving field and Ω_l is linearly proportional to the driving amplitude.

and end times respectively. Here $t_0 \equiv 0$ and $t_n \equiv \tau$. The control Hamiltonian associated with this composite pulse sequence takes the form

$$H_c(t) = \sum_{l=1}^{n} G^{(l)}(t)\frac{\Omega_l(t)}{2}\hat{\sigma}_{\phi_l} \tag{4}$$

where the function $G^{(l)}(t)$ is 1 if $t \in I_l$ and zero otherwise. The sequence of n triples $\{(\theta_l, \tau_l, \phi_l)\}_{l=1}^{n}$ completely characterizes the net effect of the applied control $(P_l = P_l(\theta_l, \Omega_l(t), \tau_l, \phi_l))$ at the end of successive pulse applications. We define the $n \times 3$ *composite pulse sequence* matrix

$$\mathbf{\Gamma}_n \quad = \quad \begin{array}{c} \\ P_1 \\ P_2 \\ \vdots \\ P_n \end{array} \begin{array}{c} \begin{array}{ccc} \theta_l & \tau_l & \phi_l \end{array} \\ \begin{bmatrix} \theta_1 & \tau_1 & \phi_1 \\ \theta_2 & \tau_2 & \phi_2 \\ \vdots & \vdots & \vdots \\ \theta_n & \tau_n & \phi_n \end{bmatrix} \end{array} \tag{5}$$

to compactly describe any arbitrary n-pulse control sequence. The entire space of such control forms therefore corresponds to an infinite set of $\mathbf{\Gamma}_n$ matrices ranging continously over all possible values taken by the control parameters. We denote this set by $\mathfrak{C}_n$ and refer to it as the *n-pulse control space*. Written formally

$$\mathfrak{C}_n := \Big\{ \mathbf{\Gamma}_n \big| \theta_l, \tau_l > 0, \;\; \phi_l \in [0, 2\pi], \;\; l \in \{1, ..., n\}, \;\; \Sigma_l^n \tau_l = \tau \Big\}$$

Noise bath model

We consider semi-classical time-dependent dephasing (detuning) and amplitude damping (relaxation) processes, captured respectively through the appearance of stochastic rotations about $\hat{\sigma}_z$ and $\hat{\sigma}_\phi := \cos(\phi)\hat{\sigma}_x + \sin(\phi)\hat{\sigma}_y$. The universal noise Hamiltonian then takes the form $H_0(t) = H_0^{(z)}(t) + H_0^{(\Omega)}(t)$ where $H_0^{(z)}(t)$ and $H_0^{(\Omega)}(t)$ denote noise interactions associated with the dephasing and amplitude noise quadratures respectively. The dephasing noise Hamiltonian is then given by

$$H_0^{(z)}(t) = \beta_z(t)\hat{\sigma}_z \tag{6}$$

where $\beta_z(t)$ is a classical stochastic process. During each pulse we also make the substitution $\Omega_l(t) \longrightarrow (\Omega_l(t) + \beta_\Omega(t)\Omega_l^{(\max)})$ where $\beta_\Omega(t)$ describes captures a (multiplicative) stochastic noise scaled by the maximum Rabi rate in a pulse segment in the amplitude of the driving field. Thus the amplitude noise Hamiltonian takes the form

$$H_0^{(\Omega)} = \beta_\Omega(t)\sum_{l=1}^{n} G^{(l)}(t)\frac{\Omega_l^{(\max)}}{2}\hat{\sigma}_{\phi_l} \tag{7}$$

and generates errors in intended rotation angle coaxial with the target rotation axis $\hat{\sigma}_{\phi_l}$.

Calculating operational fidelity & the first-order approximation

In the absence of noise, state evolution is determined by $i\dot{U}_c(t) = H_c(t)U_c(t)$ with $U_c(t)$ describing a target operation. Including the effects of noise, however, the actual evolution operator $U(t)$ satisfies $i\dot{U}(t) = (H_c(t) + H_0(t))U(t)$. A measure of the average gate fidelity $\mathcal{F}_{av}(\tau) = \frac{1}{4}\langle|\mathrm{Tr}(U_c^\dagger(\tau)U(\tau))|^2\rangle$ is then obtained using the Hilber-Schmidt inner product, effectively measuring the overlap between the intended and realized operators. These evolution dynamics however are challenging to compute due to sequential application of noncommuting, time-dependent operations giving rise to both dephasing and depolarization errors. Our approach follows the method developed by Green *et al.* [16], using average Hamiltonian theory and the generalized filter-transfer function formalism to obtain a first order approximation for gate infidelity.

Magnus Expansion

We write the total evolution operator $U(t) = U_c(t)\tilde{U}(t)$, where the *error propagator* $\tilde{U}(t)$ satisfies the Schrodinger equation $i\dot{\tilde{U}}(t) = \tilde{H}_0(t)\tilde{U}(t)$ in a frame co-rotating with the control defined by the *Toggling frame Hamiltonian* $\tilde{H}_0(t) := U_c^\dagger(t)H_0(t)U_c(t)$. Thus, in the event that $\tilde{U}(\tau) = \mathbf{I}$, the realized evolution operator $U(\tau)$ approaches the target operation $U_c(\tau)$ and errors do not affect the gate. This can be systematized by writing $\tilde{U}(\tau) = \exp[-i\Phi(\tau)]$ in terms of a time-independent effective error operator $\Phi(\tau) = \sum_{\mu=1}^{\infty}\Phi_\mu(\tau)$ with Magnus expansion terms

$$\begin{aligned}
\Phi_1(\tau) &= \int_0^\tau dt\tilde{H}_0(t) \\
\Phi_2(\tau) &= -\frac{i}{2}\int_0^\tau dt_1\int_0^{t_1} dt_2\big[\tilde{H}_0(t_1), \tilde{H}_0(t_2)\big] \\
\Phi_3(\tau) &= \frac{1}{6}\int_0^\tau dt_1\int_0^{t_1} dt_2\int_0^{t_2} dt_3\Big\{\Big[\tilde{H}_0(t_1), \big[\tilde{H}_0(t_2), \tilde{H}_0(t_3)\big]\Big] \\
&\quad + \Big[\tilde{H}_0(t_3), \big[\tilde{H}_0(t_2), \tilde{H}_0(t_1)\big]\Big]\Big\} \\
&\quad ...
\end{aligned}$$

generally taking the form of time-ordered integrals over nested commutators in $\tilde{H}_0(t)$.

One may then use vector identities (given Unitary processes) to re-express the so-called *error vector* in the toggling frame, again in an infinite power-series expansion

$$\boldsymbol{a}(\tau) = \sum_{\mu}^{\infty}\boldsymbol{a}_\mu(\tau). \tag{8}$$

This is a reexpression of the Magnus expansion error terms in the language of our control Hamiltonian. For

details of this derivation and the definition of all terms see [16].

Our theoretical predictions based on the filter-transfer function formalism involve an approximation to the *trace* or *gate fidelity* defined by

$$\mathcal{F}_{av}(\tau) = \frac{1}{4}\langle|\mathrm{Tr}(\tilde{U}(\tau))|^2\rangle = \frac{1}{4}\langle|\mathrm{Tr}(e^{-i\boldsymbol{a}(\tau)\boldsymbol{\sigma}}|^2\rangle \tag{9}$$

where $\boldsymbol{a}(\tau) = \sum_{\mu}^{\infty}\boldsymbol{a}_{\mu}(\tau)$ is the error vector given in terms of the Magnus expansion. Following the method developed in Ref. [16] and expanding the exponential in Eq. 9 we obtain

$$\mathcal{F}_{av}(\tau) = \frac{1}{2}[\langle\cos(2a)\rangle + 1] \tag{10}$$

$$= \frac{1}{2}\Big[1 + \sum_{m=0}^{\infty}(-1)^m\frac{2^{2m}}{(2m)!}\langle a^2\rangle^m\Big] \tag{11}$$

$$\langle a^2\rangle = \sum_{\mu\nu}[\langle a_1^2\rangle + \langle a_2^2\rangle + \tag{12}$$

$$+ 2(\langle\boldsymbol{a}_1\boldsymbol{a}_2^T\rangle + \langle\boldsymbol{a}_1\boldsymbol{a}_3^T\rangle + \langle\boldsymbol{a}_2\boldsymbol{a}_3^T\rangle + ...)] \tag{13}$$

with $a^2 \equiv \boldsymbol{a}(\tau)\boldsymbol{a}(\tau)^T$ the norm square of the error vector. The full expansion rapidly becomes too complex to write explicitely, however it is convenient to write $\mathcal{F}_{av} = \sum_{k=0}^{\infty}\mathcal{O}(\xi^{2k})$ where ξ is the *smallness parameter* quantifying the RMS deviation of the noise integrated over the sequence duration, $\xi \equiv \Delta\beta\tau/2$ [15]. For this series to formally converge we require $\xi^2 < 1$ (see main text and Fig. 4).

Here odd powers of ξ are omitted since these involve ensemble averages over odd powers of the noise strength and vanish under our assumption of zero-mean, Guassian-distributed random variables. Writing the $\mathcal{O}(\xi^0), \mathcal{O}(\xi^2), \mathcal{O}(\xi^4)$ classes explicitely we have

$$\mathcal{F}_{av} = 1 - \langle a_1^2\rangle - \left[\langle a_2^2\rangle + 2\langle\boldsymbol{a}_1\boldsymbol{a}_3^T\rangle - \frac{\langle a_1^4\rangle}{3}\right] + \sum_{k=3}^{\infty}\mathcal{O}(\xi^{2k}) \tag{14}$$

Immediately we see that there is a collection of terms with equal magnitude arising from *different orders* of the Magnus expansion (e.g. a_2^2 vs a_1^4). The individual terms in the series expansion of the fidelity rely on time-domain correlation and cross-correlation functions and convolution with a multidimensional control matrix capturing the effect of the control operations.

The fidelity is thus expressed explicitly in terms of noise correlations and the control matrix. For instance,

$$\begin{aligned}&\langle a_1^2\rangle \\ &= \sum_{i,j=x,y,z}\int_0^{\tau}dt_2\int_0^{\tau}dt_1\langle\beta_i(t_1)\beta_j(t_2)\rangle\boldsymbol{R}_i(t_1)\boldsymbol{R}_j^T(t_2) \\ &= \sum_{i,j,k=x,y,z}\int_0^{\tau}dt_2\int_0^{\tau}dt_1\langle\beta_i(t_1)\beta_j(t_2)\rangle R_{ik}(t_1)R_{jk}(t_2)\end{aligned} \tag{15}$$

contains all two-point noise cross-correlation functions $\langle\beta_i(t_1)\beta_j(t_2)\rangle$, for $i,j \in \{x,y,z\}$, Higher-order terms contain multipoint correlation functions (this is determined by the sum of subscript indices, as they indicate the expansion-order of the error vector).

We rewrite these terms in the frequency domain, defining the Fourier transform $\mathcal{S}_{i_1...i_n}(\omega_1,...,\omega_n)$ of an n-point cross-correlation function via

$$\langle\beta_{i_1}(t_1)\beta_{i_2}(t_2)...\beta_{i_n}(t_n)\rangle \equiv \frac{1}{(2\pi)^n}\int d\omega_1...\int d\omega_n\mathcal{S}_{i_1...i_n}(\omega_1,...,\omega_n)e^{i(\omega_1t_1+...+\omega_nt_n)} \tag{16}$$

The fidelity above can then be rewritten as

$$\mathcal{F}_{av} = 1 - \sum_{n=2}^{\infty}\left\{\frac{1}{(2\pi)^n}\sum_{i_1...i_n}\int d\omega_1...\int d\omega_n\mathcal{S}_{i_1...i_n}(\omega_1,...,\omega_n)\mathcal{R}i_1...i_n(\omega_1,...,\omega_n)\right\} \tag{17}$$

where $\mathcal{R}_{i_1...i_n}(\omega_1,...,\omega_n)$ is determined solely by the control matrix and increases in complexity at higher order. Explicit expressions for terms to arbitrary order are found in [16].

First-order fidelity approximation

Here we briefly explain the choice of fidelity metric used in the figures of the main text to produce the theory curves against which our experimental data is compared. Experimental fidelities are determined by measuring the brightness of the ion cloud after completing the control sequence, effectively yielding a projective measurement onto the $|\uparrow\rangle$ state. We denote this metric by $P_{\uparrow}(\tau) \in [\frac{1}{2}, 1]$ and refer to it as the *state fidelity*, with lower and upper bounds corresponding to complete decoherence and perfect fidelity respectively.

If the noise is sufficiently weak ($\xi^2 \ll 1$) we may truncate the series expansion for fidelity after the $\mathcal{O}(\xi^2)$ term

yielding the approximation

$$\mathcal{F}_{\mathcal{O}(1)} = 1 - \langle a_1^2 \rangle. \tag{18}$$

Here the term which dominates the measured infidelity is $\langle a_1^2 \rangle := \langle \boldsymbol{a}_1(\tau)\boldsymbol{a}_1^T(\tau)\rangle$, defined as the ensemble averaged modulus square of the first order error vector $\boldsymbol{a}_1(\tau)$. Assuming wide sense stationarity, independence and zero mean of both noise fields $\beta_z(t)$ and $\beta_\Omega(t)$ we may derive a spectral representation of $\langle a_1^2 \rangle$ of the form

$$\langle a_1^2 \rangle = \frac{1}{2\pi}\int_{-\infty}^{\infty} \frac{d\omega}{\omega^2} S_z(\omega) F_z(\omega) + \frac{1}{2\pi}\int_{-\infty}^{\infty} \frac{d\omega'}{\omega'^2} S_\Omega(\omega') F_\Omega(\omega'). \tag{19}$$

Here $S_z(\omega)$ and $S_\Omega(\omega)$ denote the dephasing and amplitude noise PSDs. The dephasing $F_z(\omega)$ and amplitude $F_\Omega(\omega)$ filter functions, on the other hand, capture the spectral response of the control sequence and are completely defined as functions of the control sequence.

As the integrated noise content increases, however, higher-order error contributions must be included; neglecting to do leads to the unphysical result that $\mathcal{F}_{\mathcal{O}(1)} \leq 0$ when $\xi^2 \geq 1$. Although computation of all higher-order contributions is challenging we may gain some insight into the full expansion by considering terms of the form $\langle a_1^{2m} \rangle \equiv \langle a_1^2 \rangle^m$ in each class $\mathcal{O}(\xi^{2m})$. This collection of terms is obtained by setting $a^2 \to a_1^2$ in Eq. 10, effectively including only the first-order *Magnus* expansion term in the expansion for *Fidelity*, yielding

$$\mathcal{F}'_{\mathcal{O}(1)}(\tau) = \frac{1}{2}\Big[1 + \sum_{m=0}^{\infty} (-1)^m \frac{2^{2m}}{(2m)!} \langle a_1^2 \rangle^m \Big] \tag{20}$$

$$= 1 - \langle a_1^2 \rangle + \frac{\langle a_1^4 \rangle}{3} - \frac{2\langle a_1^6 \rangle}{45} + ... \tag{21}$$

The oscillating sign of these terms is characteristic of the higher-order classes in converging to the true expresssion. To overcome the unphysicality of $\mathcal{F}_{\mathcal{O}(1)}$ as the noise content increases we employ a metric $\mathcal{F}_\chi$ with the physically reasonable properties that

$$\mathcal{F}_{av} \approx \mathcal{F}_{\mathcal{O}(1)} \approx \mathcal{F}_\chi, \qquad \langle a_1^2 \rangle \ll 1 \tag{22}$$

$$\mathcal{F}_{\mathcal{O}(1)} \leq \mathcal{F}'_{\mathcal{O}(1)} \leq \mathcal{F}_{av}, \mathcal{F}_\chi, \qquad \langle a_1^2 \rangle \approx 1 \tag{23}$$

$$\mathcal{F}_\chi \to 1/2 \to P_\uparrow(\tau), \qquad \langle a_1^2 \rangle \gg 1 \tag{24}$$

We may satisfy these conditions by noticing the qualitative resemblance between Eqs. 20 and 21 and the expansion for a simple exponential

$$1 - \langle a_1^2 \rangle + \langle a_1^4 \rangle - \frac{2\langle a_1^6 \rangle}{3} + ... = \frac{1}{2}\left[1 + \sum_{m=0}^{\infty} (-1)^m \frac{2^m}{m!} \langle a_1^2 \rangle^m \right] \tag{25}$$

$$= \frac{1}{2}\left[1 + \sum_{k=0}^{\infty} \frac{(-\chi(\tau))^m}{m!} \right] \tag{26}$$

where we have defined $\chi(\tau) \equiv 2\langle a_1^2 \rangle$. Hence we use the following metric in calculating fidelities to be compared with experimental data

$$\mathcal{F}_\chi = \frac{1}{2}\Big\{1 + \exp[-\chi(\tau)]\Big\} \tag{27}$$

This approximation represents the *first-order fidelity approximation*: it ignores higher-order cross correlations in the noise arising from higher-order Magnus contributions to the error vector, with diminishing overall magnitude (as given by the smallness parameter), but incorporates an approximation to higher-order terms important as the total noise-induced infidelity grows. We work in this limit throughout this manuscript.

Breakdown of the first-order fidelity approximation

As described above, the first-order fidelity ignores higher-order terms expressed as nested-integrals over cross-correlations between noise along different directions, assuming weak noise. As these contributions to gate infidelity grow in importance (for instance with α) we expect the filter-transfer-function fidelity calculations to underestimate measured error in cases where the control has filtered the noise to leading order.

We measure the probability that a π_x-pulse drives qubit population from the dark state to the bright state, $|0\rangle \to |1\rangle$, as a function of the high-frequency cutoff, ω_c, of a white dephasing bath (Fig. 4b). As the high-frequency cutoff of the noise is increased and fluctuations fast compared to the control are added to the noise power spectrum, $S_z(\omega)$, errors accumulate reducing the measured fidelity. The value of ω_c at which the fidelity drops from near unity decreases as a function of the noise strength, parametrized by α. In all cases for the primitive π pulse the fidelity calculated using the filter transfer function matches the measured data well with *no free parameters*.

Performance is notably different when studying the four-segment WAMF π-pulse, $W1$, indicated in Fig. 2d. The WAMF construction provides *first-order* filtering of time-dependent noise (red line in Fig. 2b) (effectively cancelling terms proportional to $\langle a_1^2 \rangle$), but does not provide suppression of higher-order terms in the Magnus expansion for fidelity which grow in importance with noise strength. Unlike data for the primitive gate, as the noise strength increases we observe a growing divergence between the measured fidelity and the fidelity calculated using the filter-transfer functions introduced above assuming a first order approximation (Fig. 4b).

This phenomenon is not a function of total error magnitude, but instead occurs for $\xi^2 \geq 1$ (red lines, right axis), a proxy measure indicating that we are not formally able to truncate the series expansion for fidelity at

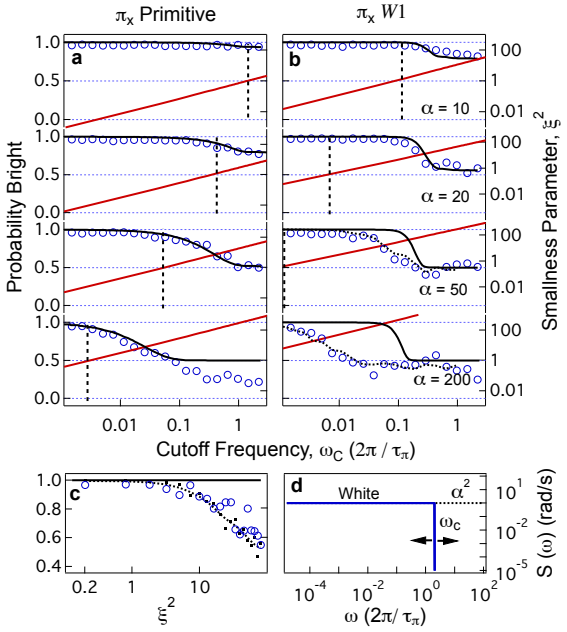

FIG. 4. Experimental validation of the breakdown of the first-order approximation to fidelity. a)-b) Measurements of operational fidelity with engineered noise for rotation $|0\rangle \rightarrow |1\rangle$ as a function of dimensionless noise cutoff frequency. Primitive rotation (a) and four-segment WAMF modulated rotation, $W1$, (b). Decay to value 0.5 corresponds to full decoherence. Each data point is the result of averaging over 50 different noise realizations. Maximum Rabi rate is fixed; filter $W1$ is conducted over time $4\tau_\pi$. Black dotted lines indicate the (smoothed) results of numerical integration of the Schroedinger equation, indicating that divergence between data and filter-function calculation is not due to experimental artefacts. Red lines (right axis) give ξ^2 employed in making the first-order filter-transfer function approximation. Vertical dashed lines indicate values of ω_c beyond which $\xi^2 \geq 1$. c) Measurements of $W1$ fidelity revealing a growing breakdown in agreement between filter-transfer function fidelity prediction and experimental data beyond $\xi^2 \gtrsim 1$, taken for $\omega_c/2\pi = 1.7\%\tau_\pi^{-1}$. d) Schematic representation of the white noise power spectrum employed.

first order and must consider higher-order error contributions [16], including Magnus terms above $\langle a_1^2 \rangle$. These measurements therefore reveal the efficacy of noise filtering and quantitatively demonstrate the bounds of the first-order fidelity approximation as breakdown routinely occurs near the predicted value $\xi^2 \geq 1$. Notably, while formal convergence of this series requires $\xi^2 \ll 1$, we find reasonable agreement between experiment and theory up to $\xi^2 \sim 5$ (Fig. 4c).

Time-domain filter order vs. Magnus order

We may formally indicate the functional dependence of the filter function on the control sequence by writing $F(\tau\omega) = F(\tau\omega; \mathbf{\Gamma}_n)$. Noise filtering (and hence error suppression) corresponds to minimizing the area under the filter transfer function in the spectral region where the noise PSD is non-negligible. We therefore define a cost function over a user-defined frequency band which may take the form

$$A(\mathbf{\Gamma}_n) := \int_{\omega_L}^{\omega_c} d\omega F(\tau\omega; \mathbf{\Gamma}_n) \tag{28}$$

to diagnose the filtering effectiveness achieved by the control sequence $\mathbf{\Gamma}_n$; the smaller the integral $A(\mathbf{\Gamma}_n)$, the more effective the noise filtering in this band. Having defined control sequences as continuous elements of the corresponding control space, for a given n we may in principle construct a variational procedure on $\mathfrak{C}_n$ to derive "values" of $\mathbf{\Gamma}_n$ satisfying a given cost function.

The filter transfer function may be approximated by a polynomial expression $F(\omega\tau) \propto (\omega\tau)^{2p}$ for some p near $\omega \approx 0$. As p increases the integral in 28, and hence the infidelity, decreases: the noise in the time domain is then said to be filtered to order $p-1$.

Equivalently stated, a control sequence $\mathbf{\Gamma}_n \in \mathfrak{C}_n$ filters time-dependent noise to order $p-1$ if $\mathbf{\Gamma}_n$ is a *concurrent zero* of the first $p-1$ coefficients in the Taylor expansion[2] of the filter transfer function about $\omega = 0$.

$$F(\omega\tau; \mathbf{\Gamma}_n) = \sum_{k=1}^{\infty} C_{2k}(\mathbf{\Gamma}_n)(\omega\tau)^{2k}. \tag{29}$$

The dependence of the expansion coefficients on our control parameters $\mathbf{\Gamma}_n$ has been made explicit, and we include only even powers of $\omega\tau$ due to the evenness of the filter transfer function. In this case $A(\mathbf{\Gamma}_n) \approx C_{2p}(\mathbf{\Gamma}_n)\frac{(\tau\omega_c)^{2p+1}}{2p+1}$ and the condition that

$$\frac{A(\mathbf{\Gamma}_n)}{C_{2p}(\mathbf{\Gamma}_n)} = \mathcal{O}\Big(\frac{(\tau\omega_c)^{2p+1}}{2p+1}\Big) \tag{30}$$

therefore implies the control sequence $\mathbf{\Gamma}_n$ filters noise to order $p-1$. This effect is visualized through the slope of the filter transfer function in the stopband on a log-log plot (Fig. 2b). A high-order filter has a higher slope in this region, indicating improved suppression of time-dependent noise.

General filter design focuses on a band of interest, permitting spectral response to diverge outside of the spectral region of interest - for instance electrical filters in the microwave may appear transparent in the THz or Hz. Therefore, in addition to the asymptotic, zero-frequency filter order $(p-1)$, we introduce a more general metric capable of describing filter performance *over an arbitrary spectral band*. The *local filter order* (p^*-1) establishes

[2] This procedure is valid for frequencies sufficiently lower than $1/\tau$ (the inverse of the total sequence duration).

	Amplitude Errors		Dephasing Errors	
	$\mu-1$	$p-1$	$\mu-1$	$p-1$
SK1	1	1	0	0
P2	2	1	0	0
B2	2	1	0	0
C1	0	0	1	1
C2 (π)	0	0	2	0
W1	0	0	1	1
W2	0	0	1	2
$\mathrm{UWMG}_{1,SK1}$	1	1	1	1

TABLE I. Comparision between suppressing *stochastic* errors to order $p-1$ (filter order) and compensating for *static offset* errors to order $\mu-1$ (Magnus order). Naming conventions for NMR sequences in top panel are consistent with the review article by Merrill and Brown [25]. The final entry corresponds to a concatenated construction described below.

that the filter-transfer function is well approximated by $F_i \propto (\omega\tau)^{2p^*}$ over the band $[\omega_L, \omega_c]$. It is this more narrowly defined metric that is used in most practical filter-design tasks, including those undertaken above.

The performance of filter-order $(p-1)$ or local filter order (p^*-1) for time-dependent noise described above must be distinguished from the order of error suppression for quasistatic errors in the Magnus expansion. The latter measure is typically used in NMR literature to the pulse sequences designed to compensate for quasistatic errors. In this regime the time dependence of the dephasing (amplitude) noise fields reduces to constants β_z (β_Ω) and the Magnus expansion terms $\Phi_\mu^{(\mathrm{DC})}$ are evaluated strictly as time integrals over *ideal* control operations scaled by powers of the offset magnitude $\beta_z{}^\mu$ ($\beta_\Omega{}^\mu$). A pulse sequence for which $\Phi_1^{(\mathrm{DC})} = ... = \Phi_{\mu-1}^{(\mathrm{DC})} = 0$ is then said to compensate offset errors to order $\mu-1$. In this case the total error operator satisfies $\Phi^{(\mathrm{DC})}(\tau) = \mathcal{O}(\Phi_\mu^{(\mathrm{DC})})$ and is dominated by the residual error proportional to the μth power in the offset magnitude.

High-order error suppression in the Magnus expansion does not imply high-order time-domain noise filtering. Table I reveals the importance of not conflating these two measures when assessing the performance of a control sequence against static vs stochastic errors. The upper panel compares the two performance measures for some well-known phase-modulated NMR sequences, the naming conventions for which are consistent with the review by Merrill and Brown[25]. For completeness, in the lower panel we also make the comparison for the novel control sequences derived in this paper.

Later we will return to the question of time-domain filter order and introduce a set of analytic design rules for filter construction based on the characteristics of our selected basis functions - the Walsh functions.

Walsh basis functions

We impose physically motivated constraints on the form of $\mathbf{\Gamma}_n$ in order to reduce the search to a manageable subspace of $\mathfrak{C}_n$, and elect to synthesize control sequences from the Walsh basis functions. The set of Walsh functions $w_k : [0,1] \to \{\pm 1\}$, $k \in \mathbb{N}$ form an orthonormal-complete family of binary-valued square waves defined on the unit interval and are the *digital analogues* of the sines and cosines in Fourier analysis. Since their formulation in the first half of the twentieth century, Walsh functions have played an important role in scientific and engineering applications. Their development and utilzation has been strongly influenced by parallel developments in digital electronics and computer science since the 1960s, with Walsh-type transforms replacing Fourier transforms in a range of engineering applications such as communication, signal processing, image processing, pattern recognition, noise filtering and so forth[35].

We summarize the relevant mathematical details of the Walsh basis, outlining two equivalent representations, *Paley ordering* and the *Hadamard representation*, both of which are useful to understand the Walsh control space.

Paley ordering

The Walsh functions are aperiodic and hence do not admit to the unique ordering according to increasing *frequency* characteristic of the sinusoids in the Fourier basis. A number of differnt orderings exist with associated definitions of the basis elements. We employ the *Paley ordering* in which basis functions are generated from products of *Rademacher functions*, a family of square waves defined by

$$R_j(x) := \mathrm{sgn}\big[\sin(2^j \pi x)\big], \qquad x \in [0,1], \qquad j \geq 0, \tag{31}$$

switching between ± 1 at rate 2^j. The Walsh function of Paley order k, denoted $\mathrm{PAL}_k(x)$, is then defined by

$$\mathrm{PAL}_k(x) = \prod_{j=1}^{m} R_j(x)^{b_j} \tag{32}$$

where $(b_m, b_{m-1}, ..., b_1)_2$ is the binary representation of k. That is, $k = b_m 2^{m-1} + b_{m-1}2^{m-2} + ... + b_1 2^0$, where $b_m \equiv 1$ is defines the most significant binary digit. Hence $R_j(x)$ is a factor of $\mathrm{PAL}_k(x)$ whenever b_j is a nonzero binary digits of k. The total number of nonzero b_j's in k define the *Hamming wieght*, denoted by r. We write $m(k)$ and $r(k)$ when it is desirable to emphasize that both m and r are functions of k. For illustration, the first 32 Walsh functions in the Paley ordering are shown in Fig. 5.

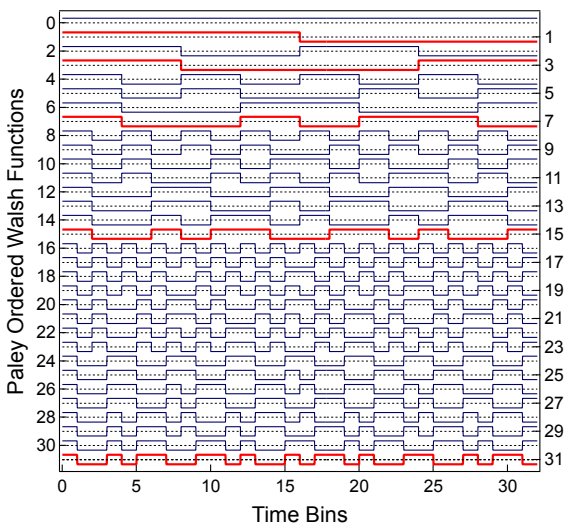

FIG. 5. The first 32 Walsh functions in Paley ordering.

Hadamard representation

For our purposes we require an expression for the piecewise-constant structure of an arbitrary superposition of Walsh functions, and therefore desire *a priori* knowledge of the locations of their various zero crossings. A general expression, however, is difficult due to the aperiodicity of the Walsh functions. It is convenient instead to use the Hadamard representation in which any continuously defined basis member $\mathrm{PAL}_k(x)$ projects completely onto a digital vector in $\mathbb{R}^{2^n}$ provided $m(k) \leq n$, which is true for the 2^n Paley orders in the set $k \in \{0, 1, ..., 2^n - 1\}$. Since these vectors have dimension 2^n and inherit the orthogonality of the $\mathrm{PAL}_k(x)$ they therefore form a *discrete* Walsh basis spanning $\mathbb{R}^{2^n}$. Such a projection is clearly possible since the fastest modulation rate in $\mathrm{PAL}_k(x)$ derives from the periodicity of $R_{m(k)}(x)$, which switches sign $2^{m(k)}$ times over $x \in [0, 1]$. The projection then involves partitioning the domain into 2^n bins and asociating the value of $\mathrm{PAL}_k(x)$ in the jth bin to the jth element $P_j^{(k)} \in \{\pm 1\}$ of the discrete digital vector

$$\boldsymbol{P}_{2^n}^{(k)} = \left[\begin{array}{ccccc} P_1^{(k)}, & P_2^{(k)}, & ... & P_{2^n}^{(k)} \end{array} \right]. \tag{33}$$

Using the so-called *Sylvester construction* [46], the 2^n-dimensional Hadamard matrix H_{2^n} is generated recursively by

$$H_{2^n} = \left[\begin{array}{cc} H_{2^{n-1}} & H_{2^{n-1}} \\ H_{2^{n-1}} & -H_{2^{n-1}} \end{array} \right] = S^{\otimes n} \tag{34}$$

$$S = \left[\begin{array}{cc} 1 & 1 \\ 1 & -1 \end{array} \right], \qquad H_1 = 1 \tag{35}$$

where S is the Sylvester matrix, and $\otimes n$ denotes $n \geq 1$ applications of the Kronecker product. In this construction $\boldsymbol{P}_{2^n}^{(k)}$ defines the $i(k) = 1 + \sum_{j=1}^{m(k)} b_j 2^{n-j}$ column (row) of H_{2^n}. The orthogonality of the Walsh basis is thereby reflected in the property that $H_{2^n} H_{2^n}^T = 2^n I$, implying the orthogonality of the Hadamard matrices.

This representation is particularly useful for efficiently constructing Walsh-synthesized waveforms. Consider an arbitrary function $f(x) = \sum_{k=0}^{N} X_k \mathrm{PAL}_k(x)$ synthesized in the Walsh basis where N sets the highest (Paley) ordered function in the construction. Then, from the above discussion, all Walsh functions in this synthesis projected onto a Hadamard matrix of dimension $\geq 2^{m(N)}$, with $\mathcal{M} = 2^{m(N)}$ giving the minimal sufficient dimension. A discrete representation of the function $f(x)$ therefore exists as a projection onto the column space of $H_{\mathcal{M}}$ by writing

$$\boldsymbol{f} = H_{\mathcal{M}} \tilde{\boldsymbol{X}}. \tag{36}$$

The column vector $\tilde{\boldsymbol{X}} = \left[\begin{array}{cccc} \tilde{X}_1, & \tilde{X}_2, & ... & \tilde{X}_{\mathcal{M}} \end{array} \right]^T$ contains the reordered Paley amplitudes X_k reoardered under the change of basis map $i(k)$ specified b

$$\tilde{X}_{i(k)} = \begin{cases} X_k & \text{for} \quad 0 \leq k \leq N \\ 0 & \text{for} \quad N < k < \mathcal{M} \end{cases}. \tag{37}$$

The vector $\boldsymbol{f} = \left[\begin{array}{cccc} f_1, & f_2, & ... & f_{\mathcal{M}} \end{array} \right]^T$ so generated then represents the piecewise constant structure of $f(x)$, with f_j giving the value taken by $f(x)$ on the jth of $\mathcal{M}$ equal subintervals partitioning $x \in [0, 1]$.

Walsh synthesis for amplitude modulated filters

Any square integrable function $f(x)$ on the interval $[0, 1]$ has a unique *spectral decomposition* in the *Walsh basis*

$$f(x) = \sum_{k=0}^{\infty} X_k w_k(x) \quad \Longleftrightarrow \quad X_k := \int_0^1 f(x) w_k(x) dx.$$

We consider a control regime referred to as *single-axis amplitude-modulation* defined by

$$\boldsymbol{\Gamma}_n \cong \{(\tau_l, \theta_l)\}_{l=1}^n, \quad \phi_l = \phi_0 \quad \forall l \in \{1, ..., n\} \tag{38}$$

where the waveform structuring the total angle swept out in each pulse segment, θ_l, is based on a linear superposition of well-defined square waves known as Walsh functions. We refer to these sequences as *Walsh amplitude modulated filters* (WAMFs), and speak of searching over the *Walsh-modulated control subspace*.

Using the above framework, we may efficiently construct arbitrary WAMFs for arbitrary pulse-segment envelopes. We define a time-varying Rabi rate $\Omega_l(t)$, $l \in \{1, ..., \mathcal{M}\}$ over the time period $t \in [t_{l-1}, t_l]$ of each segment, $l \in \{1, ..., \mathcal{M}\}$, $t \in [t_{l-1}, t_l]$. Each segment has

duration $\tau_l = \tau/\mathcal{M}$ and generates a *total rotation angle* for each segment, θ_l given by the integral

$$\theta_l = \int_{t_{l-1}}^{t_l} \Omega_l(t)dt \tag{39}$$

We now treat the rotation angles θ_l as parameters by which to optimize filter performance. For efficient filter construction, however, it is convenient to instead transform this optimization over θ_l to an optimization over a Walsh spectrum. This is achieved by writing$\theta_l = \theta_l(X_0, X_1, ..., X_N)$ with the dependence on the Walsh spectra defined by the Hadamard-matrix equation

$$\vec{\boldsymbol{\theta}} = \big(\theta_1,\ ,\theta_2, ..., \theta_{\mathcal{M}}\big)^T = (\tau/\mathcal{M})H_{\mathcal{M}}\tilde{\boldsymbol{X}}. \tag{40}$$

Defined in this way, the $\mathcal{M}$-segment arbitrary-envelope construction achieves total gate-rotation angle $\Theta = \sum_{l=1}^{n}\theta_l = X_0\tau$, completely determined by the spectral amplitude of PAL_0. All symmetry-based design rules carry over, regardless of the modulation envelope for an individual pulse segment.

Square pulses

The special case of square pulse segments is treated here as it allows a reduction in synthesis complexity and is compatible with many experimental systems. We may replace the time-dependent Rabi rate $\Omega(t)$ over a single segment with a piecewise-constant (over a single-segment) construction used in Walsh synthesis over a complete pulse sequence

$$\Omega(t) = \sum_{k=0}^{N} X_k \mathrm{PAL}_k(t/\tau), \qquad t \in [0, \tau]. \tag{41}$$

This permits synthesis over the Rabi rate per segment rather than the total rotation angle, which is often simpler in experimental settings. Substituting $\Omega(t/\tau)$ for $f(x)$ in Eq. 36 we obtain $\boldsymbol{\Omega} = H_{\mathcal{M}}\tilde{\boldsymbol{X}}$. The vector $\boldsymbol{\Omega} = \big[\ \Omega_1,\ \Omega_2,\ ...\ \Omega_{\mathcal{M}}\ \big]^T$ thus defines a sequence of modulated Rabi rates each functionally dependent on the Walsh amplitudes

$$\Omega_l = \Omega_l(X_0, X_1, ..., X_N), \qquad l \in \{1, ..., \mathcal{M}\}. \tag{42}$$

The WAMF is then defined by explicitly writing $\boldsymbol{\Omega}$ as an additional column leading the representation of Eq. 5, yielding the form

$$\boldsymbol{\Gamma}_{\mathcal{M}} = \begin{array}{c} \\ P_1 \\ P_2 \\ \vdots \\ P_{\mathcal{M}} \end{array} \begin{array}{c} \begin{array}{cccc} \Omega_l & \theta_l & \tau_l & \phi_l \end{array} \\ \begin{bmatrix} \Omega_1 & \frac{\Omega_1\tau}{\mathcal{M}} & \frac{\tau}{\mathcal{M}} & \phi_0 \\ \Omega_2 & \frac{\Omega_2\tau}{\mathcal{M}} & \frac{\tau}{\mathcal{M}} & \phi_0 \\ \vdots & \vdots & \vdots & \vdots \\ \Omega_{\mathcal{M}} & \frac{\Omega_{\mathcal{M}}\tau}{\mathcal{M}} & \frac{\tau}{\mathcal{M}} & \phi_0 \end{bmatrix} \end{array} \tag{43}$$

Here the degree of freedom associated with τ_l has apparently been removed. This reflects the fact that the choice of τ_l has been transformed into the choice of Walsh basis functions in the synthesis, each contributing its characteristic temporal profile. The remaining degrees of freedom reside in functional dependence of Ω_l on the Walsh spectrum and our variational search is thus limited to the subspace of $\mathfrak{C}_{\mathcal{M}}$ effectively spanned by $\boldsymbol{X}$.

Negative Walsh spectral amplitudes may produce negative valued Rabi rates under a linear superposition. For a pulse of the form $P_l = \exp\big[-i\Omega_l\tau_l\hat{\sigma}_{\phi_l}/2\big]$ the negative sign of Ω_l may be absorbed into the spin operator ϕ_l corresponds physically to the application of a π phase shift int he diriving field. This follows from the fact that $\hat{\sigma}_{\phi+\pi} = -\hat{\sigma}_{\phi}$, which is clear from the definition of our spin operator in Eq. 3. Thus including negative-valued Walsh spectral amplitudes generally produces single axis control only up to a sign change.

Gaussian pulse envelopes

The square-envelope pulses studied experimentally in the main text are easy to generate in our experimental system but may prove difficult in other settings where abrupt amplitude shifts at timestep-edges produce significant pulse distortion. Here we show that achieving Walsh synthesized filters using a common Gaussian pulse envelope yields comparable results with a simple re-optimization of Walsh-synthesis coefficients.

The square amplitude-modulated waveform is here replaced with a smoothly varying pulse envelope in each segment, each associated with a specific rotation angle θ_l subject to optimization. We assume a Gaussian profile $G_l(t; \mu_l, \sigma_l)$ defined on $t \in [t_{l-1}, t_l]$ with mean μ_l and standard deviation σ_l. Specifically, we construct

$$G_l(t; \mu_l, \sigma_l) = \frac{\theta_l}{C_l\sigma_l\sqrt{2\pi}} \exp\Big[-\frac{(t-\mu_l)^2}{2\sigma_l^2}\Big] \tag{44}$$

$$\mu_l = \frac{t_{l-1} + t_l}{2} \tag{45}$$

$$\sigma_l = g\tau/\mathcal{M} \tag{46}$$

with μ_l the segment midpoint and σ_l expressed as a multiple g of the segment duration. The normalizing factor

$$C_l := \int_{t_{l-1}}^{t_l} \frac{1}{\sigma_l\sqrt{2\pi}} \exp\Big[-\frac{(t-\mu_l)^2}{2\sigma_l^2}\Big]dt \tag{47}$$

is included to ensure the total rotation implemented by the Gaussian pulse in the lth segment is given by $\int_{t_l}^{t_{l-1}} G_l(t; \mu_l, \sigma_l)dt = \theta_l$. We now impose the same structure on the segment rotations θ_l as presented above in Eq. 40. Defined in this way, the $\mathcal{M}$-segment Gaussian-pulse sequence shares with the square WAMF construction the property that the total gate rotation angle

$\Theta = \sum_{l=1}^{n} \theta_l = X_0 \tau$ is completely determined by the spectral amplitude of PAL_0.

We may therefore construct a Gaussian-pulse variation on any candidate WAMF such that, having set g to some value relevant to the control hardware, the smooth pulse sequence remains strictly parametrized in the Walsh spectrum $\boldsymbol{X}$. In particular, filter optimization proceeds in the same manner as for ordinary Walsh-modulated control by minimizing the cost function with respect to the Walsh spectrum.

Analytic design rules

An advantage of Walsh synthesis is that the well-defined spectral properties and symmetries of the Walsh functions may be employed to further restrict the search space available for filter construction.

First, in practice the achievable filter order *over the entire stopband* is limited by the number of constituent control operations; one may achieve higher p at the cost of higher n. The maximum achievable value of p for a given filter is set by the power-law expansion of the filter for the single Walsh function with the highest Paley order for a given n. As has been shown previously, all Walsh functions with given Hamming weight of the Paley order have the same power-law expansion near zero frequency [36]. Therefore, in principle, every doubling of n increases the maximum achievable time-domain filter order by one. The Walsh functions highlighted in red in Fig 2c and Fig. 5 represent those with the highest Paley-order Hamming weight for a given n. Nonetheless we find that in general we are able to construct filters with higher order than prescribed *over narrow regions in the stopband*, as a result of Walsh synthesis (see the multiple slopes for the blue line in Fig 2b).

In filter construction we may further constrain the form of a candidate pulse sequence by imposing required physical properties on the sequence, such as fixing the total rotation angle of the Bloch vector in order to implement a target logic operation. In order to proceed we then partition the Walsh spectrum $\boldsymbol{X} \equiv (\boldsymbol{X}_\nu, \boldsymbol{X}_\rho)$ into spectral amplitude classes $\boldsymbol{X}_\nu$ and $\boldsymbol{X}_\rho$ to be treated as *variational* and *fixed* parameters respectively. Fixed parameters set the physical state transformation of interest while the remaining unconstrained components in $\boldsymbol{X}_\nu$ serve as tuning parameters by which to minimize $A(\boldsymbol{X}_\nu; \boldsymbol{X}_\rho)$.

The primary constraint in WAMF constructions is that the total rotation angle executed depends only on the value of X_0, the zeroth order spectral component; it sets the effective average Rabi rate for the WAMF. This can be seen as follows. First observe all Walsh functions of higher than zeroth order are *balanced* in the sense that $\int_0^1 \mathrm{PAL}_k(x)dx = \delta_{0k}$. For the control field defined by Eq. 41 the total gate rotation angle $\Theta = \int_0^\tau \Omega(t)dt$ then takes the form

$$\begin{aligned}\Theta &= \int_0^\tau \sum_{k=1}^{N} X_k \mathrm{PAL}_k(t/\tau)dt \\ &= \tau \sum_{k=1}^{N} X_k \int_0^1 \mathrm{PAL}_k(x)dx \\ &= \tau \sum_{k=1}^{N} X_k \delta_{0k} = X_0 \tau.\end{aligned}$$

In this case the net gate rotation $\theta = \Theta \bmod 2\pi$ is given by

$$\theta = X_0 \tau \bmod 2\pi \tag{48}$$

implying the necessary constraint on X_0 in order to achieve a desired θ.

Next, we observe that the Walsh functions have distinct parity, but that filter constructions mandate symmetric constructions in order to enact a target operation and provide effective noise cancellation. The result is that odd-parity Walsh functions may generally be excluded from the variational search. While this is not necessarily strictly required (multiple odd parity Walsh functions may in principle be added with opposite signs to produce net symmetric constructions), it is convenient and effective to restrict the synthesis space to the so-called CAL subset of the Walsh functions.

Our reduced search problem may then be represented formally by replacing $\boldsymbol{\Gamma}_\mathcal{M} \to (\boldsymbol{X}_\nu, \boldsymbol{X}_\rho)$ in Eq. 28 to obtain

$$A(\boldsymbol{X}_\nu; \boldsymbol{X}_\rho) := \int_0^{\omega_c} d\omega F(\tau\omega; \boldsymbol{X}_\nu, \boldsymbol{X}_\rho) \tag{49}$$

with the variational search now restricted to the subspace spanned by $\boldsymbol{X}_\nu$ with the $\boldsymbol{X}_\rho$ held constant.

First order WAMFs

As a first application of the above result, we derive a family of nontrivial gates decoupled to first order against dephasing noise by constructing a pulse sequence from the synthesis $\theta(t) = \frac{\tau}{4}\Big(X_0 \mathrm{PAL}_0(t/\tau) + X_3 \mathrm{PAL}_3(t/\tau)\Big)$. Note that $\theta(t)$ is only formally defined at the end of pulse segments. That is, we set $\boldsymbol{X}_\rho \equiv X_0$ and $\boldsymbol{X}_\nu \equiv X_3$. In this case $N = 3$ and $\mathcal{M} = 4$, so the minimal pulse duration is $\tau/4$. Using Eq. 37 we obtain $\tilde{\boldsymbol{X}} = \left[\, \tilde{X}_1,\ \tilde{X}_2,\ \tilde{X}_3,\ \tilde{X}_4 \,\right]^T = \left[\, X_0,\ 0,\ 0,\ X_3 \,\right]^T$, yielding the minimal Hadamard representation

$$\boldsymbol{\theta} = \frac{\tau}{4} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{bmatrix} \begin{bmatrix} X_0 \\ 0 \\ 0 \\ X_3 \end{bmatrix} = \frac{\tau}{4} \begin{bmatrix} X_0 + X_3 \\ X_0 - X_3 \\ X_0 - X_3 \\ X_0 + X_3 \end{bmatrix} \tag{50}$$

These sequences therefore span the control subspace parametrized by ($\tau = 1$)

$$\text{WAMF } \mathcal{O}(1) \quad = \quad \begin{array}{c} \\ P_1 \\ P_2 \\ P_3 \\ P_4 \end{array} \begin{array}{c} \begin{array}{cccc} \Omega_l & \theta_l & \tau_l & \phi_l \end{array} \\ \begin{bmatrix} X_+ & \frac{X_+}{4} & \frac{1}{4} & 0 \\ X_- & \frac{X_-}{4} & \frac{1}{4} & 0 \\ X_- & \frac{X_-}{4} & \frac{1}{4} & 0 \\ X_+ & \frac{X_+}{4} & \frac{1}{4} & 0 \end{bmatrix} \end{array} \tag{51}$$

where we have defined $X_\pm = X_0 \pm X_3$. This choice is motivated by the fact that Paley order $k = 3$ corresponds to the lowest order non-constant Walsh function with even symmetry about $\tau/2$ (Fig. 5). Hence 50 is the simplest Walsh modulated form which includes the zeroth order Walsh function in the synthesis and which possesses time-reversal symmetry about the sequence midpoint. The former property ensures a nontrivial gate angle is executed. The latter is chosen due to the observation in dynamic decoupling literature that using time-symmetric building blocks often improves the performance of the sequence compared to sequences formed by time-asymmetric building blocks[47][43]. In this construction we have maintained a leading column of rabi rates, Ω_l, as would be appropriate for the square-pulse forms used in the main text.

Gaussian pulse construction

Adding to the results presented in the main text, constructing the 4-segment filter $W1$ using square pulse segments, we examine the Gaussian-pulse variation here. The cost function $A_z(X_3; X_0) = \int_{\omega_L}^{\omega_c} d\omega F_z(\omega\tau; X_3; X_0)$ may be computed by partitioning the time domain into a large number N_s of subintervals on which the continuous Gaussian envelope is treated as approximately constant. Fig. 6 shows a two-dimensional representation of $A_z(X_3; X_0)$ integrated over the interval $\omega \in [10^{-9}, 10^{-6}]\tau^{-1}$, with $g = 1/6$ and $N_s = 100$. The value of $\text{Log}_{10}\big[A_z(X_3; X_0)\big]$ is indicated by the color scale. Total sequence length is normalized to $\tau = 1$ in this data, so the total gate rotation angle $\Theta \equiv X_0$ is given directly by the X_0-axis. Regions in blue represent effective (first-order) filter constructions, where the cost function is minimized.

We conclude useful filter construction using Gaussian pulses is a simple matter of re-optimization in the Walsh-synthesis framework. This is readily achieved using a Nelder-Mead optimization of $A_z(X_3; X_0)$ for any particular choice of $g, \omega_L, \omega_c, X_0$ or N_s, in a manner precisely the same as for square envelopes.

Universal Filters by Concatenation

Phase-modulated sequences robust against amplitude noise may also be found in the Walsh basis, yielding *Walsh phase-modulated filters* (WPMFs) analogous to WAMFs, and implementing arbitrary target rotations θ. There are a variety of techniques to construct such WPMFs, but we use analytic design rules in which a target rotation is performed (with some error due to noise), and phase-modulated segments are added in order to produce the net filtering effect. The simplest WPMF adds two segments phase modulated according to Walsh function PAL_1 with coefficient X_1, subject to the constraint that one enacts the desired driven rotation by θ.

The result of this approach yields a WPMF that is identical to the NMR sequence SK1, with value $X_1 = \cos^{-1}(-\theta/4\pi) \equiv \phi_{\text{SK1}}(\theta)$ [25], as represented

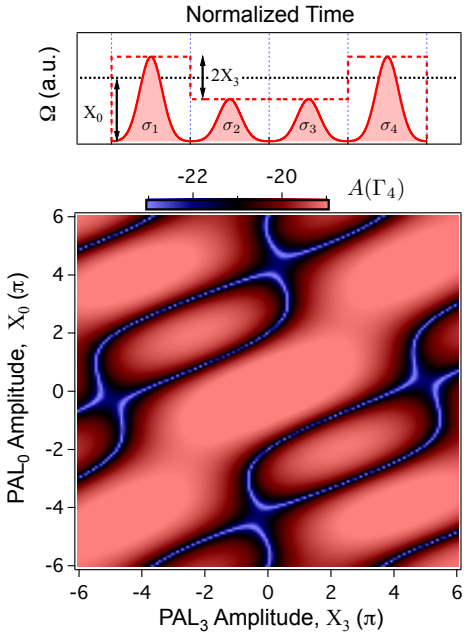

FIG. 6. Construction of the first-order Walsh amplitude modulated dephasing-suppressing filter using Gaussian-shaped pulse segments. a) Schematic representation of Walsh synthesis for a four segment gate of discrete Gaussian segments. Walsh synthesis determines the overall amplitude of Gaussian pulses with fixed duration and standard deviation, setting the effective pulse area in each segment. b) Two-dimensional representation of the integral metric defining our target cost function, $A(\mathbf{\Gamma}_4)$ integrated over the stopband $\omega \in [10^{-9}, 10^{-6}]\tau^{-1}$. Areas in blue minimize $A(\mathbf{\Gamma}_4)$, representing effective filter constructions. The X_0 determines the net rotation enacted in a gate while X_3 determines the modulation depth, as represented in a).

$$\mathbf{\Gamma}_3^{(\mathrm{SK1})} = \begin{array}{c} \\ P_1 \\ P_2 \\ P_3 \end{array} \begin{array}{c} \begin{array}{cccc} \Omega_l & \theta_l & \tau_l & \phi_l \end{array} \\ \begin{bmatrix} \Omega_0 & \theta & \tau_\theta & 0 \\ \Omega_0 & 2\pi & \tau_{2\pi} & \phi_{\mathrm{SK1}} \\ \Omega_0 & 2\pi & \tau_{2\pi} & -\phi_{\mathrm{SK1}} \end{bmatrix} \end{array} \tag{52}$$

$$\Omega_0 = \frac{\theta + 4\pi}{\tau}, \qquad \tau_\theta = \frac{\theta}{\Omega_0}, \qquad \phi_{\mathrm{SK1}}(\theta) := \cos^{-1}\Big(-\frac{\theta}{4\pi}\Big). \tag{53}$$

Note that the Walsh timing construction only holds in the two correction steps represented above. Following a similar route allows one to construct a sequence with modulation given by PAL_3 which is formally identical to the three-segment (four-timestep) phase modulation given by gate $P2$ [25].

These WPMF sequences perform as first-order time-dependent noise filters, captured in the form of $F_\Omega(\tau\omega)$, and noted in Table I (column 2). For example, filter functions for the WPMF that is equivalent to SK1 are shown in Fig. 7b, revealing first-order filtering of amplitude noise, but not dephasing noise.

We may now concatenate WAMFs and WPMFs in order to simultaneously filter universal noise. We focus on an explicit example providing first-order amplitude and dephasing noise filtering. The basic procedure is to implement each constant-amplitude segment of a four-segment WAMF, $W1$, using a constant-amplitude phase-modulated sequence robust against amplitude noise. As a reminder, the noise-filtering performance of $W1$ is shown in Fig. 7c. Here we use the WPMF≡SK1 sequence for the phase modulation. We refer to the concatenated gate as a Universal Walsh Modulated Gate, $\mathrm{UWMF}_{1,SK1}$.

Referring to Eq. 51, the WAMF filter is similarly written $P_3(X_+/4,0)P_2(X_-/2,0)P_1(X_+/4,0)$. Concatenation then involves the operator substitutions

$$P_1(X_+/4,0) \to \mathrm{SK1}^{(1)}(X_+/4) \tag{54}$$

$$P_2(X_-/2,0) \to \mathrm{SK1}^{(2)}(X_-/2) \tag{55}$$

$$P_3(X_+/4,0) \to \mathrm{SK1}^{(3)}(X_+/4). \tag{56}$$

The composite structure for $\mathrm{UWMF}_{1,SK1}$ is shown in Fig. 7a. Here the SK1 phase flips $\phi = \pm\phi_{\mathrm{SK1}}$ within each segment of the WAMF profile are indicated by the oppositely oriented hatching; $\phi = 0$ is indicated by white fill. The dephasing and amplitude filter functions for the concatenated sequence are shown in Fig. 7d, indicating effective filtering of both amplitude and dephasing noise.

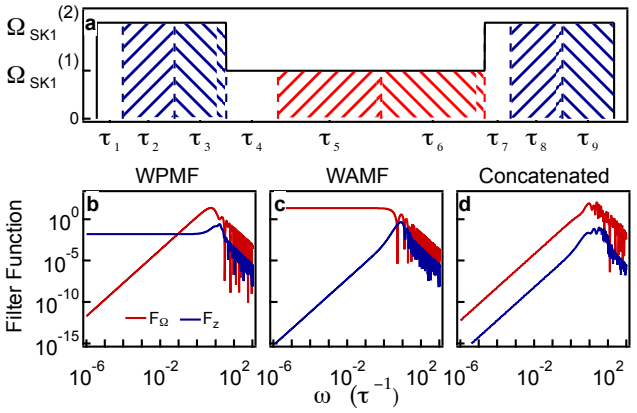

FIG. 7. Concatenation scheme for universal noise suppression. a) Concatenation of WPMF≡SK1 within first order WAMF sequence yielding $\mathrm{UWMF}_{1,SK1}$. White fill indicates rotations enacted with $\phi = 0$; orientation of hatching denotes SK1 phase flips $\phi = \pm\phi_{\mathrm{SK1}}$. b) Filter functions for WPMF≡SK1 sequence. c) Filter functions for four-segment WAMF sequence, $W1$. d) Filter functions for concatenated sequence.

Ytterbium Ion Trapping

We use trapped $^{171}\mathrm{Yb}^+$ ions as our experimental platform; a detailed description of related experimental approaches appears in [26, 48]. A linear Paul trap enclosed in an ultra-high vacuum (UHV) chamber is used to trap several hundred $^{171}\mathrm{Yb}^+$ ions as a small homogeneous ensemble (in magnetic field and microwave field amplitude). Doppler cooling of the ions is achieved using 369 nm laser light, slightly red-detuned from the $^2\mathrm{S}_{1/2}$ to $^2\mathrm{P}_{1/2}$ transition. Additional lasers near 935 nm and 638 nm are employed to depopulate metastable states.

Our qubit is the 12.6 GHz hyperfine splitting between the $^2\mathrm{S}_{1/2}\,|F=0,m_f=0\rangle$ and $^2\mathrm{S}_{1/2}\,|F=1,m_f=0\rangle$ states. For notational simplicity we will designate $|0\rangle$ and $|1\rangle$ to these states respectively. Addition of a 2.1 GHz sideband to the 369nm laser using an electro-optic modulator permits high-fidelity state preparation in $|0\rangle$. For details of ion loading, laser cooling, state preparation, and state detection see [26]. While we typically employ a small ensemble of ions, the system behaves similarly to single-ion experiments in our lab, and benefits from both high-fidelity state initialization and projective measurement - the system does not bear similarity to NMR-style ensembles.

State detection is achieved by counting 369 nm photons scattered from the ions and converting to a probability that the Bloch vector lies at a particular location along a meridian of the Bloch sphere. This measurement is susceptible ion loss in the ensemble and both laser amplitude and frequency drifts over long timescales, resulting in variable maximum and dark count rates over time. We therefore employ a normalization and Bayesian estimation procedure for state detection, see [26].

An important advantage of this system is that the selected qubit transition is first order insensitive to magnetic field fluctuations; the measured free-evolution in our setup is $T_2 \approx 4$ s, limited by coherence between

the qubit and the master oscillator [26]. Coherent rotations between the measurement basis states are driven by using the magnetic field component of resonant microwave radiation. The Rabi rate for driven oscillations reaches ~ 14 μs in our system, with typical operation near ~ 50 μs. Rotations are implemented about an axis $\vec{r}$ lying on the xy-plane of the Bloch sphere and set by the phase of the microwaves as $\vec{r} = (\cos\phi(t), \sin\phi(t), 0)$. Driven operations, characterized by randomized benchmarking, exhibit a mean fidelity in excess of 99.99%.

Noise Engineering

In the laboratory we rely on engineering noise in our control system to provide a method to accurately reproduce decoherence processes of interest. We begin with a desired noise power spectral density in either the amplitude or detuning quadrature (or both), assuming they are statistically independent. From this power spectrum, defined by the noise strength α, the exponent of the power-law scaling p, the comb spacing ω_0, and the high-frequency cutoff $\omega_c \geq J\omega_0$, we numerically generate time-domain vectors for amplitude and frequency errors. Noise is injected into the system by adding these modulation patterns on top of the control sequence being implemented (e.g. a pulse of radiation for implementing a π-pulse)using IQ modulation in our vector signal generator [26].

Randomized Benchmarking

We use randomized benchmarking as a tool for resolving small gate errors which cannot be resolved in the application of a single gate. Our randomized benchmarking sequence consists of interleaved $\pi/2$ and π pulses each applied along axes randomly selected from $\pm x$ and $\pm y$. Each sequential pair of $\pi/2$ and π rotations is referred to a computational gate. A given randomized benchmarking sequence consists of l computational gates followed by a final correcting gate which is selected such that the aggregate Unitary operation applied is a π rotation. For each l we measure 50 randomizations (dots in Fig. 3f), and in each randomization average over 20 different realizations of a white dephasing noise bath. Each realization, in turn, employs 20 measurements in our Bayesian state-detection algorithm, in addition to associated normalization experiments.

Comparisons of $W1$ to primitive π rotation performance in randomized benchmarking is conducted via replacement of all π pulses with $W1$ constructions, again about randomly selected axes. In either case the $\pi/2$ rotation is achieved using a primitive gate, although we have also validated that replacement of the $\pi/2$ gates with WAMF constructions yields net improvement in gate fidelity.